Fermilab

Computing Sector

Eighteen Month Progress Report

"Run II Data Preservation Project"

Version 1.0

10/10/2014

CD-DocDB #5186

PREPARED BY:

Robert D. Kennedy, Joseph B. Boyd, Kenneth Richard Herner, Bodhitha Jayatilaka, Willis Sakumoto, Stephen P. White

CONCURRENCES:

Rob Roser	Date
Project Sponsor	

Table of Contents

1. Project Overview

The purpose of this project is to implement data preservation for the Tevatron Run II experiments. The Run II data sets are very unique data samples of 1 TeV protons colliding with 1 TeV anti-protons, and represent a considerable investment. It is unlikely another data sample like this will ever be collected again. This sample is in many ways complimentary to the LHC data sample of multi-TeV protons on multi-TeV protons. A number of important "legacy" Tevatron physics measurements have been identified. While the majority of existing analysis tools will be available and supported for five years after end of data taking, the datasets will have value in generating and checking physics results throughout another decade.

Data preservation means different things to different people. The approach for this project will be to:

- curate the data from CDF and D0 Run II (both simulated and actual) by continuously migrating the data to modern storage media,
- maintain the full infrastructure ability to generate new Monte Carlo samples, simulate them, reconstruct them, and process them, and
- maintain the ability to perform physics analysis on both the simulated and actual data.

With this, a complete physics analysis can be carried out on Run II data well into the future.

2. Project Status

This report covers the progress as of **June 30, 2014** on the Run II Data Preservation Project.

2.1. Project Scope

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project scope. There have been no scope changes since that charter version. Note that the project scope for each Run II experiment varies a bit according to the needs of each individual experiment.

2.2. Project Schedule

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project time frame. The current project timeline is:

- Jan – Jun 2013: Evaluate current state and plan the R2DP system
  - July 2013: Assess 6 Month milestones
- July – Dec 2013: Implement the R2DP pilot system, address some open issues
  - First significant changes/deployments permitted by experiments in Aug/Sep 2013
  - Jan 2014: Assess 12 Month milestones
- Jan – Jun 2014: Gap-fit the pilot system, address all open issues, finalize R2DP system plan
  - July 2014: Assess Internal 18 Month milestones
- July – Dec 2014: Deploy the R2DP system for both experiments and close-out the project
  - Nov-Dec 2014: Assess 24 Month milestones and begin project close-out

Due to moderate project underspending and later-than-expected delivery of IF job submission tools, the project is will not finish in December 2014 as was originally planned. While the production R2DP system itself will be released and in operation by December 2014, the project will continue to assist in training and support of the early adopters of the Run II Data Preservation system for a brief time into 2015.

The high-level project Gantt Chart is shown in Figure 1 below. Detailed R2DP tasks are tracked on the project's SharePoint site (https://sharepoint.fnal.gov/project/DataPreservation/default.aspx) in the "Data Preservation tasks" list.

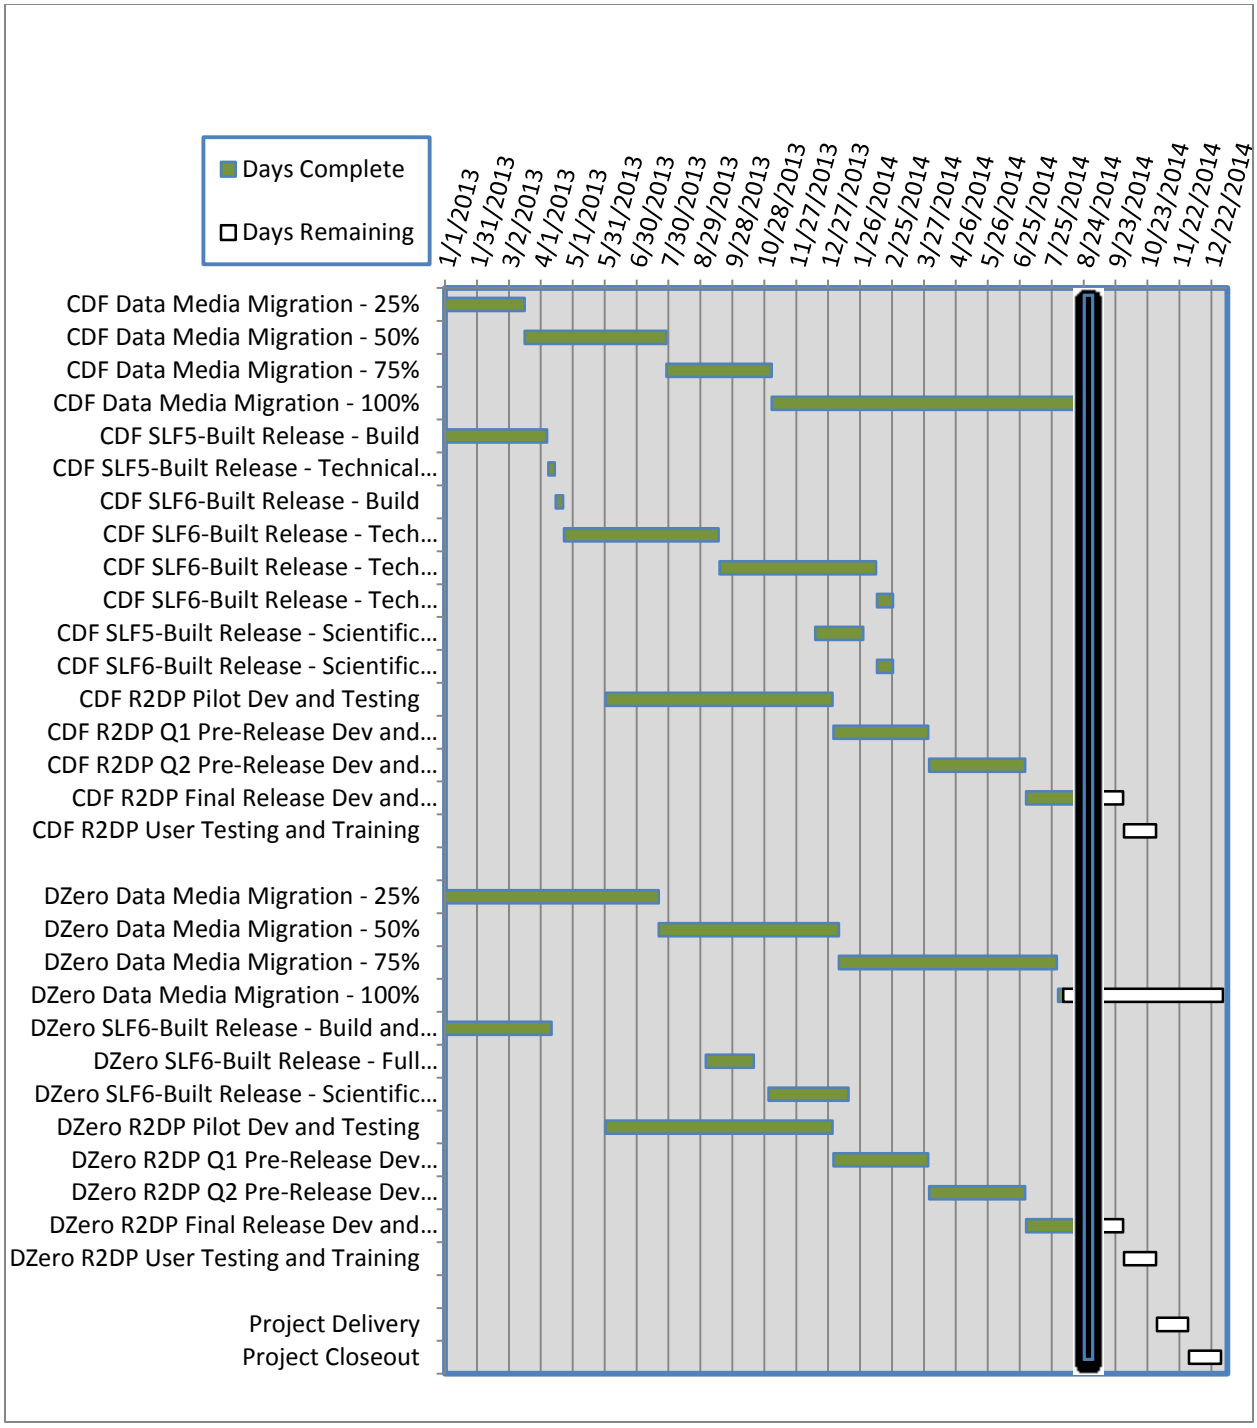

Figure 1: Timeline for Identified Run II Data Preservation Work with progress (green) as of 9/1/2014

The data migration goals were nominally achieved for the 18-month mark, at least 75% of data was migrated.

The goal of the CDF Offline Release is to produce a frozen offline release natively built on SL6, with an SL5 build as an intermediate deliverable. The SL5 build was validated and used in the CDF R2DP pilot (achieved 12-month project goal). The SL6 build planned to be ready in February 2014 was delayed by the investigation and fix of obscure crashes in simulation and was released at the end of March 2014. This was the version used for the R2DP initial release as well as the adaptation to and testing of CVMFS use and SAMWeb use (achieved 18-month goal).

The goal of the DZero Offline Release is to produce a frozen offline release built on SL6 using some SL5-compatibility libraries. This release was validated in January 2014 and used for the DZero R2DP pilot (achieved 12-month project goal). This build was used for the R2DP initial release as well as the adaptation to and testing of SAM data file caching using dCache (achieved 18-month goal).

2.3. Project Resources

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project resources and organization at a high level. Willis Sakumoto joined the Project Core Team in May 2014 to transition into the role of the R2DP leader for CDF beginning June 1, 2014. At the end of June 2014, Bo Jayatilaka left the project team. Bo and Willis worked together to successfully achieve a smooth and effective role transition. Otherwise, the Project Core Team remains the same as is defined in the project charter. The extended project team continues to evolve as Run II experiment support transitions to more efficient smaller-scale operations, which has led to more opportunities to align data preservation “operations” work with this project’s work.

2.4. Project Cost

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project budget. There have been no changes to the budget since that charter version.

The project budget for FY13 + FY14 is 6.97 FTE-years in SWF, and zero dollars in M&S. As of June 30, 2014, the project had expended 4.23 FTE-years of effort, continuing to be underspent in the reporting.

While we planned to engage more resources to report their project effort to the project activity code, there was confusion with this request as Run 2 Data Preservation is also a title used for reporting operations work. Even if effort had been reported to the project, however, only Bo’s and Ken’s effort are mapped to the R2DP project task codes behind the scenes. So, other contributors’ effort would still have been charged to Run 2 Data Preservation operations task codes even if we had succeeded in getting more project effort reported to the project activity code.

With Willis’s different employee status relative to Bo and Ken, future project cost reports will report on the project task code spending as well as the effort reported. This is not particularly useful at this stage as Willis only joined the project towards the end of this reporting period.

The cumulative project effort to date for FY13 and FY14 are shown in Figures 2 and 3 below.

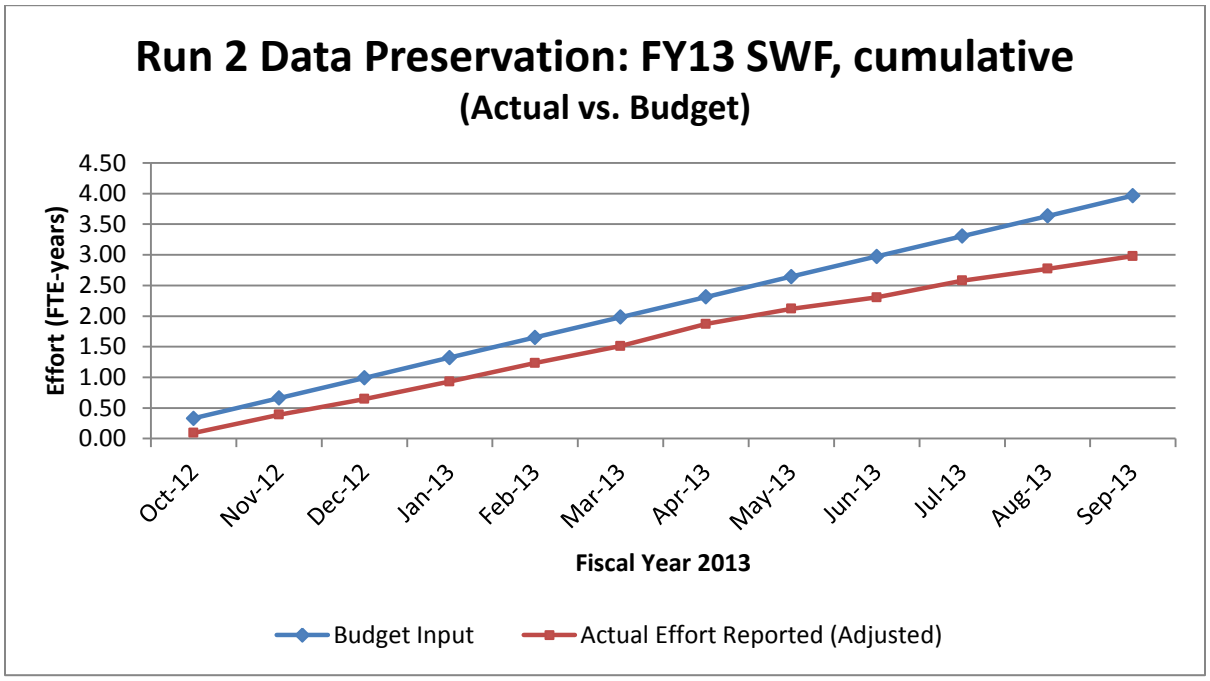

Figure 2: Cumulative Effort Reported in FY13 for the Run II Data Preservation Project

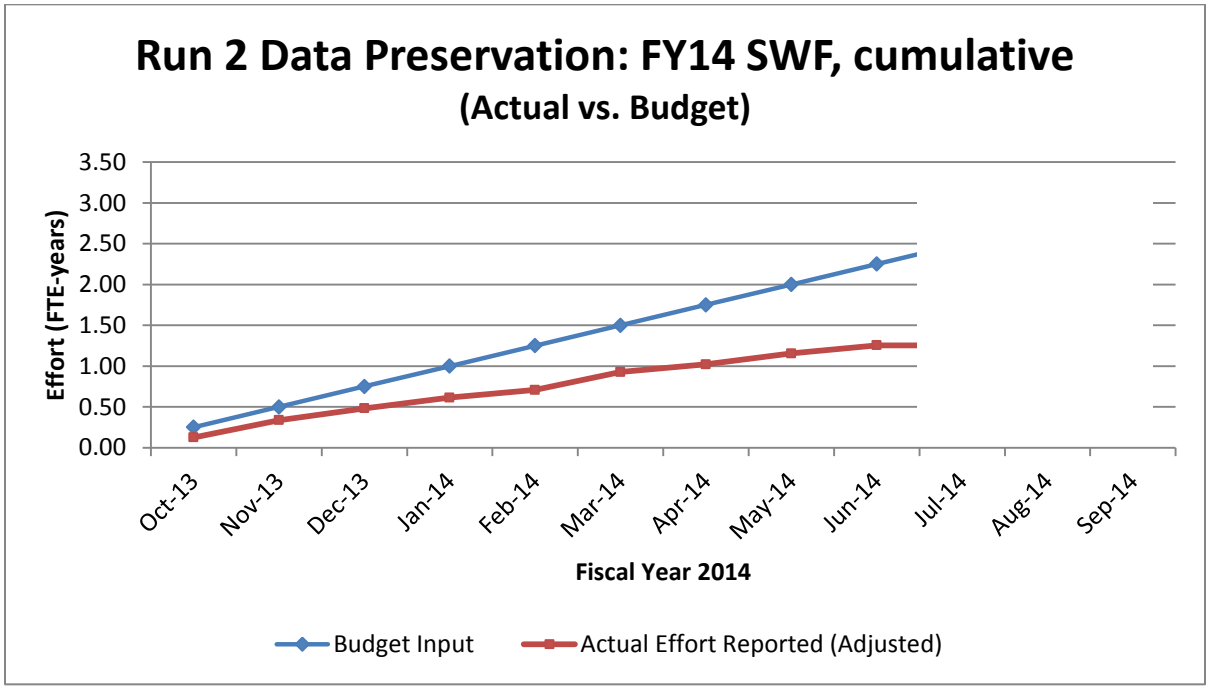

Figure 3: Cumulative Effort Reported in FY14 as of June 30, 2014 for the Run II Data Preservation Project

3. Progress on 18 Month Milestones

The project charter does not define any 18 Month Milestones. The project core team, however, created an ambitious set of 18-Month Milestones in order to track progress in the critical period leading up to the release of the Run 2 Data Preservation systems and the close of the project.

3.1. Data Migration

Goal: Complete 75% of data migration to modern T10K tape technology

Description: The project oversees this request to the Mass Storage Services (MSS) group, who will actually perform the actual work. The project and the MSS group work with the experiments to assess priorities in data migration, to insure that on-going data analysis takes priority over this data migration. In order to achieve the 25% data migration goal per each 6 months of the project, each experiment data migration effort should sustain on average about 1% of the data migration per week.

Results: as of 6/30/2014

- CDF: 99.9% of data migrated
  - The CDF data migration has essentially been completed. There are 3 tapes remaining to be closed. After 6/30/2014, an error in configuration allowed 3 more tapes to be written to old media, requiring subsequent migration to T10k media, for a total of 6 tapes to be addressed to complete the CDF data tape media migration.
- DZero: 78.3% of data migrated (based on "Migrated value)
  - The DZero data migration sustained a rate sufficient to migrate all data by Dec 31, 2015. However as the files and file families have become smaller on average, a larger fraction of tapes are being left unclosed by the migration since they may be used to store later small file families. The number of tapes "Closed" that has been used to mark progress in the migration (now 70.8%) no longer reflects the volume of data migrated to new media, so we are choosing to switch to the number of tapes "Migrated" (now 78.3%) to judge progress for the migration.

Assessment: **Goal achieved.**

Next Steps: Project should continue to track data migration rates to ensure continued success.

3.2. Complete Transition to Long-Term Infrastructure

Goal: Complete transition to long-term infrastructure: shared services and VM's wherever possible

Description: The project has decided to move Run II services/systems to shared services and virtual machines wherever possible. (See Open Issues 8f, 9a, and 9b.)

Results: The main production servers for the R2DP system have either been virtualized or will be virtualized when the physical servers are next refreshed. Major experiment-specific services have been moved to shared services to the extent that all parties agree makes sense. Some systems, however, would require far greater effort to move to a fully-shared, non-experiment-specific configuration than first thought. Sharing Mass Storage Systems, for example, would require combining PNFS namespaces which is technically challenging, and would yield little benefit as use of the systems continues to fall.

Assessment: **Goal achieved.**

Next Steps: None.

3.3. Address Most of the Open Issues

Goal: Address most of the Open Issues

Description: The project has identified Open Issues and a plan to resolve them. Open Issues are high-level unresolved issues worthy of reporting to the Project Sponsor. Some Open Issues are in fact risk mitigation believed to be necessary to ensure that the Run II Data Preservation system is sustainable by the Fermilab Scientific Computing Division with minimal effort from the Run II experiments.

Results: Appendix 1 contains the updated Open Issues list. Of the identified 11 high-level Open Issues:

- 7 are fully CLOSED.
- 4 are partly closed, with some planned work remaining.

Assessment: **Goal achieved.**

Next Steps: The project will complete the remaining Open Issues. The remaining work is mostly related to the completion of the R2DP system deployment, making it the default.

3.4. Test A Friendly Analysis Using Future System

Goal: Complete and test the R2DP system for at least one friendly analysis for each experiment

Description: In order to test some of the design of the Run II Data Preservation system once the basic system is completed, the project will perform at least one “friendly” analysis for each experiment to demonstrate that scientific analysis can be done in the new system, and also that the validation procedure can be performed in the new system.

Results:

- Ther are still a few tasks to complete the R2DP system before an analysis as ultimately intended (unlike other parts of the progress report, the status described here is as of early Oct. 2014):
- CDF:
    - CDFSOFT2 6.2 releases suitable for testing by the collaboration are available, and are to be released for collaboration wide tests and acceptance in Oct. 2014. Once approved, they become the CDF frozen legacy releases for /cdf/code and CDF CVMFS systems.
    - The IF jobsub client and tools feature set is being upgraded to accommodate all required CDF CafSubmit features. As of Oct. 2014, this upgrade is close to full completion. With this upgrade. there is a one-to-one feature/option correspondence between CDF CafSubmit and IF jobsub. The CafSubmit like wrapper over IF jobsub is trivial.
- DZero:
    - Improve how user results are copied back to the user in bulk. This is implemented in the forthcoming IFDH release expected in Oct. 2014, and will require trivial changes to DZero production scripts to take advantage of.

Assessment: **Goal partly achieved, remainder is a work in progress.**

Next Steps: The project is confident that the remaining work will be completed and a friendly analysis performed before the end of the project. At this point, analyses have already been run for both experiments as a part of testing the R2DP system, but the finishing touches to the R2DP system require outside parties to complete work on the IF job submission tools.

3.5. Design the Validation Procedure

Goal: Design the validation procedure based on existing procedures

Description: In order to validate that the R2DP system is working, especially in a low-use environment, a simple procedure will be defined with a well-defined expected result.

Results:

- CDF: the validation procedure is defined and will become operational for the R2DP system once cdfsoft2 and job submission are ready.
- DZero: the validation procedure is defined, and is being implemented and tested as part of the overall R2DP system testing.

Assessment: **Goal achieved.**

Next Steps: The project will continue to build and test the implementation of the validation procedure as part of the overall R2DP system testing.

3.6. Assemble Final Documentation

Goal: Assemble first version of the final documentation set.

Description: Each experiment requires sufficient documentation to enable users to make use of the R2DP system in the context of offline software analysis for that experiment.

Results:

- CDF: The documentation is the same, as CDF CVMFS and IF jobsub are under wrapper layers that make them appear as /cdf/code and CafSubmit. Direct use of CDF CVMFS only involves sourcing a setup script, and after this, all other CDFSOFT2 setups are the same.
- DZero: The first version of this documentation is in writing and has been presented to the experiment.

Assessment: **Goal partly achieved, remainder is a work in progress.**

Next Steps: The project will complete the documentation once the final R2DP system is fully deployed and a friendly analysis completed to identify the documentation changes required.

3.7. Identify Training Approach

Goal: Identify who and how experimenter training will be performed.

Description: Working with the experiments, the project will identify how R2DP training will be delivered.

Results:

- CDF: Update web documentation, and rely on the R2DP system working like the past system. No specific user training planned.
- DZero: Ken will train physics group conveners, and the conveners will then be responsible for training delivery in their area of scientific analysis responsibility.

Assessment: **Goal achieved.**

Next Steps: Update documentation and follow-up as needed with experiments' physics conveners.

3.8. Business Model and Risk Plan for R2DP System

Goal: Outline the business model and risk plan for the R2DP system

Description: The business model and risk plan for the R2DP system is intended to address questions related to the long-term maintenance of the system, such as "who will pay for tape media if another tape media migration is required several years from now" or "how will we address a security patch to Oracle RDBMS that breaks the user interface in experiment software". This was done at a high level in the early part of the project across several documents, and the stretch goal here is to do so in more detail in one document.

Results:

- The business model and risk plan have already been documented in general terms.
- No additional work though has been done to enhance the documentation of the business model and risk plan for the R2DP system. We will address this in more detail once the implementation of the R2DP system is completed.

Assessment: **Goal partly achieved, remainder is deferred.**

Next Steps: Write a more detailed description of the business model and collect together the risk plans to create a more complete and detailed document.

4. Progress Summary

Overall, the R2DP project continues to achieve its goals. Tape data media migration continues at an acceptable pace. Though CDF's frozen offline release has taken longer to validate than expected, the timeline of the project has not been affected yet since a SL5 validated release was available. The CDF Offline database upgrade was a success, and Frontier has moved to virtual machines with a long-term support plan. DZero RD2P leveraged work by DASPOS resource to adapt to CVMFS use. Progress was made setting up a DZero dCache to allow testing of dCache underneath SAM at DZero. Planning has begun for a more detailed feature delivery schedule as both experiment's efforts have achieved success in the initial pilot. We expect this success to continue as the project works towards delivering a tested initial release of R2DP system for both experiments by July 1, 2014.

Appendix 1. Open Issues List

We have grouped the Open Issues into two main categories, issues from the users' point of view and issues from the service providers' point, to reflect our goal of delivering a usable and sustainable system. Documentation references are to the R2DP Project site and require Fermilab SERVICES login.

To date, 7 of 11 Open Issues are closed and 4 of 11 Open Issues are partly closed: issues 6,7,8,11.

How will data analysis be performed in the Run II Data Preservation system? – *users' point of view*

1. How are we going to perform bookkeeping?
    a. Plan: Both Run II experiments will move from using SAM to using SAMWeb for a sustainable data bookkeeping solution.
    b. CDF Documentation: SAMWeb for CDF (Illingworth).pdf
    c. DZero Documentation: SAMWeb for DZero.pdf
    d. Status: CLOSED (18-month)
    e. Timeline: Project has closed this question, though the implementation work continues.
        i. SAMWeb has been integrated and tested in both experiments.
        ii. CDF: the cdfsoft2 6.2 cut of software will use SAMWeb by default. This is expected to be released around 9/15/2014.
        iii. DZero: There is concern about making changes before the end of the 5-year period (post-TeV shutdown). Ken may bring this up after upcoming production release, to turn SAMWeb on by default.

2. How will a user be authorized to access data?
    a. Answer: The R2DP system will continue with the existing experiment management plan for conceptual authorization even after the Run II experiments end.
        i. This depends on some informal Run II collaboration existence that can define who is a member of the "experiment" and thus qualified to access the data.
        ii. The technical means of authorization will remain the same as well.
    b. Status: CLOSED (6-month)

3. How will jobs be run? How will the user start his/her job?
    a. Plan: Submit jobs using the Intensity Frontier (IF) job submission system.
        i. DZero Plan: We plan to integrate existing DZero job submission scripts with the IF system to hide this change from the end-users.
        ii. CDF Plan: We plan to use the same CAFsubmit. Nothing changes for end-users.
    b. Plan: Jobs will run on a grid or cloud node. Details of this are under discussion.
    c. Status: CLOSED (18-month)
    d. Timeline: Project has closed this question, though the implementation work continues. Project will complete the re-implementation of the main experiment job submission interfaces using the IF job submission tools as they become available.
        i. As of this writing (10/1/2014), all needed features appear to be in the development version of IF job submission tools.
        ii. Move of these features to production is scheduled for late September (ITIL change request made).
        iii. Estimate for full implementation and deployment is Dec 31, 2014.

4. Is our MC framework flexible enough to interface with new generators?

a. Answer: Yes, because the input format to the MC framework can always be flat ASCII text file.
b. Status: CLOSED (6-month)

5. Can we use new data analysis programs as they become available?
    a. Answer: This is up to the user to test and implement, not the project. Past experience with ROOT encourages us to believe that users will be able to use new versions of ROOT on existing data either directly or after applying a ROOT-supplied conversion tool. Note:
        i. New versions of ROOT will not be centrally managed or deployed by project.
        ii. Experiment code will not be built and validated against new versions of ROOT by the project.
    b. Status: CLOSED (6-month)

6. How will users be supported for data analysis questions?
    a. Answer: While not strictly in project scope, perceived success of system depends on this. The elements the project may pursue to encourage this outcome include the following:
        i. Identify and address gaps in support process due to fewer active users and analyses. This is an Experiment task. The Project is a stakeholder, but not responsible for execution.
        ii. Update user documentation known to be out-of-date. This is an Experiment task. The Project is a stakeholder, but not responsible for execution.
        iii. User analysis support system – there should be reduced expectations for support response time, and greater reliance on documentation. This is an Experiment task. The Project is a stakeholder, but not responsible for execution.
    b. Plan: The Project will help identify areas where data analysis support may be weak to the experiment as we enable example analyses to run in the Run II Data Preservation system, to ensure the perceived quality of the project deliverable is not reduced by lack of analysis support.
    c. Status: OPEN – Work in progress, expect completion at 24-month mark.
    d. Timeline:
        i. CDF: documentation for running high-level processes is a work in progress.
            1. Example: running the simulation.
            2. User interface for job submission does not change (the reason for the jobsub work).
            3. Saving to tape: Ray is working on this task within the high-level process
        ii. DZero: documentation of the changes in the R2DP system is already done.
            1. Final docs are simple: how to get started with new system, add on to the documentation that is already there.
            2. May add suggestions for new infrastructure like the copy-back task (24 months)

How will the infrastructure be maintained for a Post-2015 system? – *service provider's point of view*

7. What needs to be upgraded to current versions to re-establish version currency?
    a. Data migration to modern T10k media
        i. Plan: track data migration. Work with the Mass Storage group and the experiments to address bottlenecks in the migration.

ii. Status: OPEN (24-month goal)
    1. Met 75% migration goal by July 1, 2014.
    2. On target to complete migration by Dec 31, 2014.
iii. Migration Progress Tracking:

https://sharepoint.fnal.gov/project/DataPreservation/SitePages/TapeMigration.aspx

b. CDF: Offline Oracle RDBMS to 11
    i. Plan: upgrade the Offline development RDBMS first and test operation. Then, after the summer 2013 conferences, upgrade the Offline production RDBMS.
    ii. Update: Upgrade was successfully performed 9/10/2013.
    iii. Status: CLOSED (12-month).
c. CDF: Frontier/Squid
    i. Plan: upgrade to the most modern version of Frontier and its underlying toolset that are usable by CDF.
    ii. Update: All servers updated in July 2013. Support is transitioning from CDF experimenter to SCD (with experimenter consulting) per agreement.
    iii. CDF Documentation: CDF Frontier.docx
    iv. Status: CLOSED (12-month).
d. Overall Status: OPEN
e. Timeline: Project will close the least topics of this question by the end of the project.

8. What should be changed/upgraded to reduce the risk of obsolescence in existing infrastructure?
    a. Convert to using new SAM server
        i. Plan: Both Run II experiments will move from using SAM to using SAMWeb for a sustainable data bookkeeping solution. See discussion for Open Question 1.
            1. Both experiments have enabled the use of SAMWeb. Neither experiment has yet made SAMWeb the default over the old SAM server solution.
        ii. CDF Documentation: SAMWeb for CDF (Illingworth).pdf
            1. The current legacy CDFSOFT2 6.2 releases uses either the old SAM/Corba or SAMWeb interfaces on CDFGrid and jobsub/FNAL-GP. A user settable environment switch sets the access mode.
        iii. DZero Documentation: SAMWeb for DZero.pdf
        iv. Status: Open (24-month goal)
    b. DZero: Calib DB Servers
        i. Plan: Deploy a virtual node with the new SLF6-based releases of all Calibration DB servers. This will replace the Calibration DB servers currently on d0dbsrv20. We plan to execute this move after the Summer 2013 conference season.
        ii. Note: This only partially overlaps Issue (9a) since the service software is also being updated and ported to work on a newer operating system.
        iii. Update 2/2014: One Calib DB Server upgraded and deployed. The remainder are ready to be deployed at next regularly scheduled downtime.
        iv. Status: Closed (6-month).
    c. Migrate logbooks to supported tools
        i. CDF Status: already using the supported tool.
        ii. DZero Status: Migration has been converted to ECL.
            1. See http://dbweb0.fnal.gov/ECL/dzero
        iii. Status: Closed (6-month).
    d. Migrate some auxiliary databases, database applications to supported tools

i. Most auxiliary databases have been treated. The exceptions are noted below:
ii. CDF:
    1. The People DB does not have a plan yet, but is not a high priority and is arguable whether it is in project scope.
    2. Internal Notes going to Spires. This is not in project scope, but it being tracked as part of the overall experiment R2DP system.
iii. DZero:
    1. Whod0, Speakers Bureau will be dumped into a simple archive format. These are part of the DZero central web server. DZero is not concerned about this however.
    2. We may want a risk mitigation, like dump it to a large HTML/XML table.
    3. Internal Notes already migrated to Spires.
iv. Status: Open (24-month goal)

e. DZero: Transition from CORBA to HTML
    i. Plan: Transition plan documented, but experiment chose not to proceed
    ii. Documentation: DZero Documentation: D0 Calibration DbServers, D0_CORBA_Emails, CORBAtoHTTP_Effort_Summary
    iii. Agreement: We met with the DB group and the D0 algorithms group and agreed that the html transition will occur if and when CORBA stops working. The CORBA infrastructure will not be fixed. D0 understands that the CORBA-to-HTTP transition may take longer to complete if the work is done later rather than now owing to experts potentially moving on to other projects. D0 also understands that MC production will be unavailable for the entire CORBA-to-HTML transition period. D0 accepts the risks of not performing the transition now in order to save valuable manpower for other projects during the 2013-14 timeframe.
    iv. Status: Closed (6-month).
f. How long are certain production servers needed, before being retired?
    i. Plan: We have a comprehensive list of candidate production servers. Many have already been migrated as driven by the refresh cycle or risk issues.
    ii. Main production servers have either been virtualized or will be virtualized when the physical servers are next refreshed. In a few cases, all parties agree to run less-used niche-specific production servers until they die.
    iii. Ray Culbertson has migrated the tape uploader infrastructure to the supported CDF data server machines: fcdfdas11 and 12. The validation test was the uploading of 3800 physics ntuple files. All but one file were successfully written to tape, which Ray verified to be T10K [(tag)(library) = CDF-10KCF1]. We consider this task to be completed. The case of one or two files not uploading automatically is normal, and the failed-to-upload files are fixed/uploaded by hand.
    iv. Documentation: DPPlan.docx
    v. Status: Closed (18-month)
g. Overall Status: CLOSED (18-month)

9. How/where will we adapt to enterprise architecture guidance to reduce operations costs?
    a. Move services on “physical nodes” to virtual nodes (tied to Issue 8f)
        i. Main production servers have either been virtualized or will be virtualized when the physical servers are next refreshed.
        ii. In a few cases, all parties agree to run less-used niche-specific production servers until they die, such as the CDF upload servers.

iii. Status: Closed (18-month)
iv. Documentation: DPPlan.docx

b. Move services on experiment-specific systems to shared systems where possible
   i. This has been done to the extent that all parties agree makes sense. Some systems, such as the Mass Storage Systems, would require far greater effort to move to a non-experiment-specific configuration than it would yield benefits. In this example, combing PNFS namespaces would be technically challenging.
   ii. Status: Closed (18-month)

c. Overall Status: CLOSED (18-month)

10. How is this system going to be technically sustainable from 2015 through 2020? Roadmap?
    a. SLF6 versus SLF7
       i. Plan: Focus on an all-SLF6 system that operates in a virtual environment. The host CPUs may run a different OS, but the analysis environment will be SLF6.
       ii. Documentation: D0 s/w beyond 2020.pdf, CDF task force report, DPPlan.docx
       iii. Status: Closed (6-month).
    b. Oracle RDBMS
       i. Plan: Investigate the potential roadmap for Oracle. Develop a risk plan that takes into account the cost of upgrading versus not-upgrading especially given recent positive experience operating on unsupported Oracle RDBMS versions at CDF.
       ii. Documentation: DPPlan.docx
       iii. Status: Closed (6-month).
    c. DZero: SAM Cache and dCache (for local data caching)
       i. Plan: Transition to using dCache for this role, a product with long-term support.
       ii. Documentation: Dzero Job Submission Talk
       iii. Status: Closed (18-month).
    d. Overall Status: CLOSED (18-month)

11. How is this system going to be financially sustainable from 2015 through 2020? Support plan?
    a. Determine how to prove the R2DP system is operational after changes.
       i. Plan: Identify and document official validation test procedures, standard analyses
       ii. Plan: Determine if automating some or all validation testing is feasible.
       iii. Status: Open. – WORK IN PROGRESS. (24 month)
          1. CDF test suite is waiting on the cdfsoft2 6.2 release. Exists.
          2. DZero test suites: B Physics has their own code base. Conveners not sure of value given rate of codebase divergence of studies after the validation. Top validation is done and in use. WMass will not be ready on this timescale.
    b. CDF: Frontier/Squid Support
       i. Plan: Develop a plan to migrate support responsibility from CDF to SCD.
       ii. Status: Squid support has not yet been formally worked out with SCD.
       iii. Status: Closed (12-month).
    c. How will costs for a possible future media migration be addressed?
       i. Plan: engage SCD management on this question. This is heavily dependent on future CMS choices of media, equipment, and robotics.
       ii. Status: Open. – Document in Close-out and Seek Customer Sign-Off
    d. Overall Status: OPEN
    e. Timeline: Project will address these topics by the end of the project.